arXiv:1702.01703v2 [q-bio.GN] 8 Feb 2017

Variant tolerant read mapping using min-hashing

Jens Quedenfeld[1,3] and Sven Rahmann[2,3]

[1] Chair of Theoretical Computer Science, Technical University of Munich, Germany
jens.quedenfeld@in.tum.de
[2] Genome Informatics, Institute of Human Genetics, University Hospital Essen, University of Duisburg-Essen, Essen, Germany
Sven.Rahmann@uni-due.de
[3] Bioinformatics, Computer Science XI, TU Dortmund, Germany

Abstract. DNA read mapping is a ubiquitous task in bioinformatics, and many tools have been developed to solve the read mapping problem. However, there are two trends that are changing the landscape of readmapping: First, new sequencing technologies provide very long reads with high error rates (up to 15%). Second, many genetic variants in the population are known, so the reference genome is not considered as a single string over ACGT, but as a complex object containing these variants. Most existing read mappers do not handle these new circumstances appropriately.
We introduce a new read mapper prototype called VATRAM that considers variants. It is based on *Min-Hashing* of q-gram sets of reference genome windows. Min-Hashing is one form of *locality sensitive hashing*. The variants are directly inserted into VATRAMs index which leads to a fast mapping process. Our results show that VATRAM achieves better precision and recall than state-of-the-art read mappers like BWA under certain cirumstances. VATRAM is open source and can be accessed at `https://bitbucket.org/Quedenfeld/vatram-src/`.

1 Introduction

In bioinformatics, DNA read mapping has become an important basic step for many sequencing analysis tasks. Given millions of short DNA fragments (so called *reads*) over the alphabet $\Sigma = \{A, C, G, T\}$ and a reference genome which is a long DNA string that is many magnitudes longer than the reads, the problem is to locate these reads in that genome. A typical (short) read length is 100–300 nucleotides [14], the human reference genome has a length of approximately 3 billion nucleotides. Since not all individuals are identical, there are some differences between a given read and the corresponding interval on the reference genome. Furthermore, the sequencing machines are not perfect and produces sequencing errors, therefore, the best match is searched for, e.g. according to the edit distance.

To solve the read mapping problem efficiently, usually an index data structure is used to find an exact match between a short substring of the read and the

reference genome. Afterwards, the exact match (called *seed*) is extended to a full alignment. Many popular read mappers (such as BWA [12] or Bowtie2 [10]) use the FM index which is based on suffix arrays and the Burrows Wheeler Transformation [7]. However, if there are many differences between the read and the reference genome, then there are few unique long seeds that can be found, but many unspecific short seeds, so common read mappers have difficulties to map reads with many errors efficiently

This problem becomes more and more important for two reasons. First, there are new sequencing technologies (for example from Pacific Biosciences [6]) which produce long reads with up to 60 000 nucleotides and a high error rate of 10%–15%. Second, thousands of individual human genomes have been sequenced in the last years, so many frequent variants are known. It is commonly accepted that the human genome is not represented well by a single string over the alphabet Σ. To improve the read mapping process it is necessary to consider these known variants. About 90% of the differences between the individuals of one species are *single nucleotide polymorphisms* (SNPs), which are substitutions of one single nucleotide in DNA. Other important variant types are insertions or deletions of one or more nucleotides as well as large structural rearrangements. Index data structures based on the FM index have problems to handle these new circumstances, so new types of indexes have to be explored.

One such alternative is *locality sensitive hashing* (LSH) on q-gram sets of reference intervals. LSH has been already used for finding similarities between different DNA sequences [4] as well as for genome assembly [1] of PacBio reads. However, as far as we know LSH was not used for read mapping yet. We have developed a prototype of a read mapper called VATRAM (*VAriant Tolerant ReAd Mapper*) which is able to consider known variants. VATRAM uses *MinHashing* [2] which is one specific form of LSH.

This paper is organized as follows: In section 2 we describe in detail how the index of VATRAM works and how it is created. Afterwards (section 3) we explain how reads can be found in the reference genome using this index and how they are aligned using our variant tolerant aligner. Section 4 contains several experiments where VATRAM is compared to other read mappers.

Preliminary ideas, implementations and experiments were reported in an internal report [9] and a Master's thesis [16]; the present article contains our summarized findings. The source code of VATRAM is available at `https://bitbucket.org/Quedenfeld/vatram-src/`.

2 Index creation

2.1 Basic idea

The index of VATRAM is created as follows. First, each chromosome is divided into windows of length w, such that the windows are slightly longer than the typical read length n. The distance between two consecutive windows is denoted by $o \leq w$, so the windows may (and usually do) overlap. A typical configuration is $w = 1.4n$ and $o = 1.25n$ [16].

For each window, the set of contained q-grams (substrings of length q) is determined. If the window contains a SNP which appears with a known population frequency higher than δ, then the q-grams containing the SNP are added to the window's q-gram set. A typical value is $\delta = 0.2$.

If there are two or more SNPs with a distance lower than q, then we add all combinations of q-grams to the set. However, if are too many SNPs within a given substring of length q, than these SNPs are ignored, because adding all combinations would enlarge the q-gram set too much. In the extreme case there would be q consecutive positions where all four bases (A,C,G,T) are allowed. So there would be 4^q possible q-grams for that substring. If we added these to the window's q-gram set, it would contain all possible q-grams, so there is no information which could help us to map a read to this window. By default a limit of $l = 3$ q-grams for each q-gram position is used.

Each q-gram set is mapped to a single value using a technique called min-hashing [2]. For that purpose we conceptually choose an arbitrary permutation of all q-grams uniformly at random. A q-gram set is mapped to the smallest q-gram according to the order defined by this permutation. The resulting q-gram is called the *signature value* of the window. If we compare two q-gram sets Q and Q' the min hash property holds:

Lemma 1 (Min-hash property). *Let $\mathcal{Q}$ be the set of all strings over Σ of length q. Given two sets $Q \subset \mathcal{Q}$, $Q' \subset \mathcal{Q}$ and the set Π of all permutations on $\mathcal{Q}$, let $\pi \in \Pi$ be a random permutation. For any $Q \subset \mathcal{Q}$, define $h(Q) := \min_\pi(Q)$. The probability that Q and Q' are hashed to the same value $h(Q) = h(Q')$ is equal to the Jaccard coefficient of Q and Q',*

$$P(h(Q) = h(Q')) = \frac{|Q \cap Q'|}{|Q \cup Q'|}. \tag{1}$$

A proof can be found in [2].

Of course, in practice it is not possible to do this, because there are $(4^q)!$ different permutations, so $\log_2(4^q!)$ bits are required to represent a permutation. For a common value like $q = 16$, this is about $4 \cdot 10^{10}$ bits per permutation, which is not practical.

Therefore, instead of a permutation we choose a random 32 bit word which is called *permutation value*. The q-grams are also represented by 32 bit hash values[4] The signature value is calculated using the *exclusive or* operation (XOR) between each q-gram hash and the permutation value π and then taking the minimum:

$$h_\pi(Q) = \min\{x \oplus \pi \mid x \in Q\}. \tag{2}$$

Using 32 random bits and the XOR technique instead of a true random permutation means that the pre-conditions of Lemma 1 do not hold and the

[4] If the q-gram length is larger than 16, some q-grams are mapped to the same hash value. However, we found out that longer words (e.g. 64 bit) only increases the memory consumption, but the mapping quality stays almost the same.

min-hash property may be violated [3]. However, empirical studies have shown that in practice the XOR technique approximates the desired property well [5].

Mapping each window to a single signature value is not enough to find a read. Therefore s different permutation values are used. The parameter s is called *signature length*.

2.2 Data structure

For each permutation value a data structure, such as a simple hash table, is needed to map the calculated signature values (of each window) to the particular window in the reference genome. Instead of a simple hash table, we use a two layer succinct rank data structure, because it needs less memory.

An one layer rank data structure consists of a bit array B, an offset array C and a data array D. The bit array is divided into blocks of length λ. For each block there is one entry in the offset array that indicates the number of 1's in B up to this block. The data array contains one data entry for each 1 in B, so the length of D is equal to the number of 1's in B. Given an index i in B with $B[i] = 1$, the corresponding data entry $D[j]$ can be efficiently accessed by using the offset array C. Additional information about rank data structures can be found for example in [15].

This one layer rank data structure can be directly used for our purpose. Each entry in B represents a possible signature value, so the length of B is 2^{32}. The elements of the data array are also arrays that contain the particular window references.

This data structure has two disadvantages resulting in a very high memory consumption: First, the array B needs $\frac{2^{32}}{8}$ Bytes = 512 MB space. Note that this space is needed for each permutation value, so the memory consumption of B must be multiplied with s, so for $s = 36$ we get a memory usage of already 18 GB for the bit arrays. Second, the most signature values are unique, i.e. the most arrays in D contain only one window reference. This leads to a high memory overhead.

The first problem is solved by using a two layer rank data structure which is visualized in figure 1. For the human genome with 3 billion nucleotides, there are 24 million windows (when using the default window distance $o = 125$). Thus, only 0.6%[5] of the entries are 1's and with $\lambda = 32$ at least 82%[6] of the blocks contain only zeros. The two layer rank data structure uses the arrays B^+ and C^+ in the super layer and B^-, C^- and D^- in the second layer. Note that the data array D^+ of the super layer equals the arrays B^- and C^-. The array B^- contains only those blocks of B that contain at least one "1". The super layer is needed to decide if a given block k in B is empty (i.e. $B^+[k] = 0$) or not (i.e. $B^+[k] = 1$). The data array D^- is exactly the same as before, so $D^- = D$. By

[5] Because $\frac{24 \cdot 10^6}{2^{32}} \approx 0.006$

[6] There are 24 million windows, so at most 24 million blocks can contain at least one "1". Thus, the ratio of blocks that contains only zeros is $1 - \frac{24 \cdot 10^6}{2^{32}/\lambda} = 0.8211...$

default, the block size of the super layer is $\lambda^+ = 64$ and the block size of the second layer is $\lambda^- = 32$.

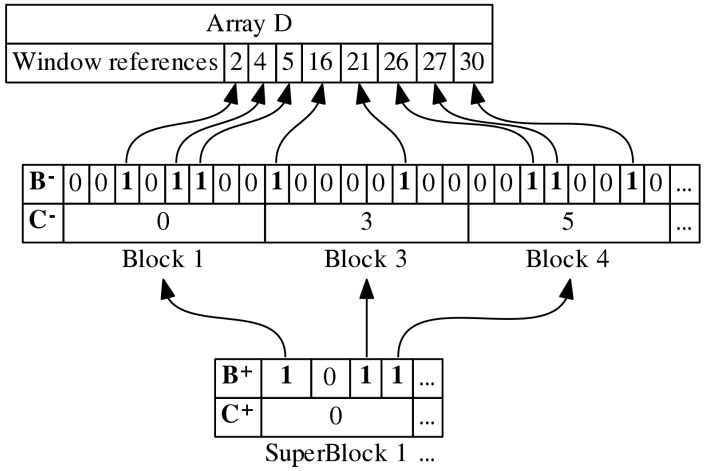

Fig. 1. Visualization of the two layer rank data structure. The figure was adapted from [9], some variable names were changes.

The second problem can be solved as follows. Usually a window reference is stored as a 32 bit value. However, for reasonable configurations there are always less than 2^{30} (about one billion) windows, so the two most significant bits can be used for meta information. If there is only one window for a given signature value, then the two bits are set to 00 and the window reference is stored directly in D^-. For signature values having exactly two corresponding windows, there is *one* additional array whose elements consists of two window references. The corresponding entry in D^- contains the index in that array and the meta bits are set to 01. Furthermore there is one array whose elements consists of up to four window references and another one whose elements are dynamic arrays. The described data structure is called *Multi Window Manager* and visualized in figure 2.

Using more permutations leads to better results in the sense that we are better able to distinguish true similarities from random hits; however, the memory consumption increases linearly with the number of permutations. For the human genome and reads of length $n = 100$ usually $s = 36$ permutation values are used which leads to a memory consumption of about 14 GB [16]. For longer reads, the window distance relative to the read length (o/n) can be decreased, so there are less windows and therefore less value pairs that have to be stored in the rank data structure, so the number of permutations can be increased.

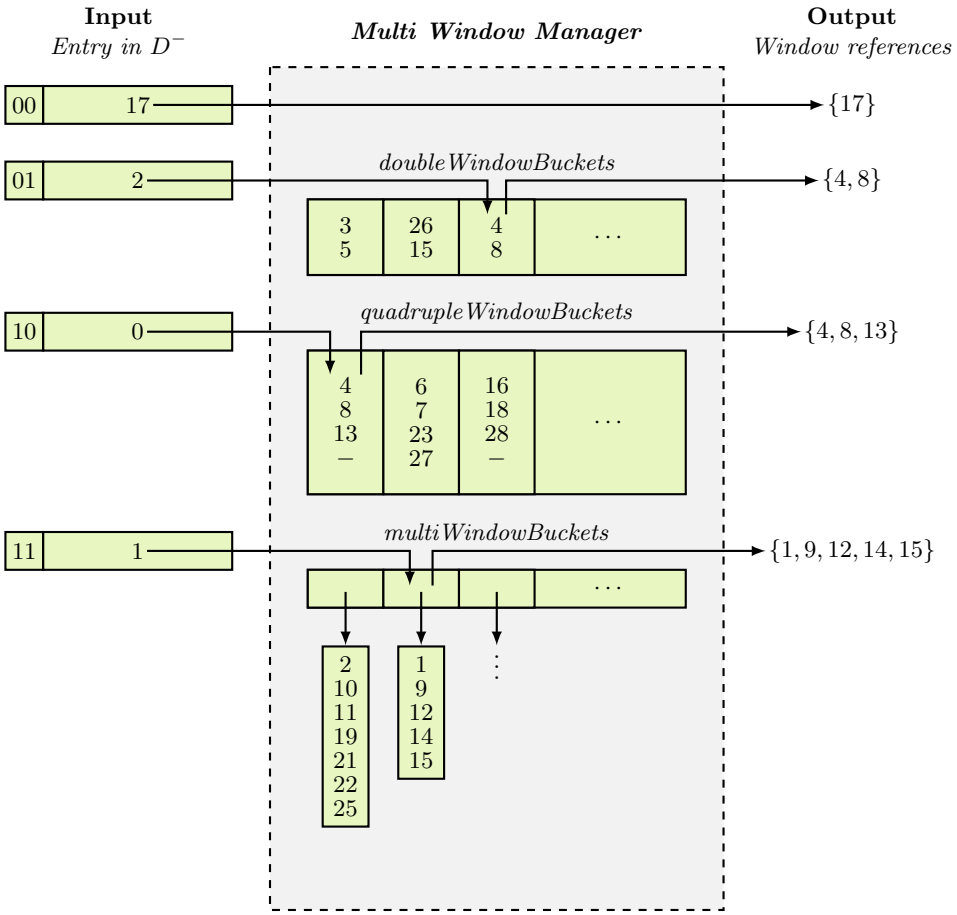

Fig. 2. Visualization of the *MultiWindowManager*

3 Read mapping

Mapping and aligning a DNA read is done in two phases. In the mapping phase, candidate windows of the reference genome are determined and converted to candidate intervals. In the alignment phase, a variant tolerant alignment between the read and the candidate intervals is performed to obtain the best alignment, taking all variants into account.

3.1 Finding reference intervals

To find a given read in the reference genome, first its q-grams set is determined. Then the signature values are determined analogous to the index creation using the same permutation values. Each signature value is looked up in the index data structure, so for each permutation value we receive a (possibly empty) list of window indices. The same procedure is done with the reverse complement of the read.

Now we count how many times each window index was found. This is done by sorting all found window indices. To accelerate this, we sort each window list in the index data structure already during the index creation. Then we use merge sort to merge the already sorted window lists.

If a read is located between two windows, there are common q-grams with both windows and therefore we probably obtain several common signature values for both windows. For this reason neighboring windows are summarized to a so-called *window sequence*. Single windows (that have no neighbors) are also represented by a singleton window sequence. Each window sequence is scored using the number of hits for each contained window. The score is calculated as follows.

Given a window sequence Ω that consists of $|\Omega|$ windows. Let $c_1, \ldots, c_{|\Omega|}$ be the number of hits for each contained window. If $|\Omega| = 1$, then the score is simply c_1. For longer window sequences the score $C(\Omega)$ is defined by

$$C(\Omega) \coloneqq \max_{i \in \{1,2,\ldots,|\Omega|-1\}} (c_i + c_{i+1}) \tag{3}$$

The idea is that a read cannot intersect with more than two windows if $n + w < 2o$ holds, which is fulfilled in the standard configuration. Therefore, adding all c_i of a long window sequence (i.e. $\sum_{i=1}^{|\Omega|} c_i$) would result in too large scores if there is a repetitive region in the reference region.

The best scoring windows sequences are selected (by default at most $\kappa = 64$). Then each window sequence is converted to an interval on the reference genome. Window sequences containing one window are enlarged depending on how often the window index was found (a window with only a few hits it is further enlarged than a window with many hits). Window sequences containing two windows are contracted, because it is likely that the read is contained in both windows. The more hits for the windows in the window sequence are found, the smaller is the resulting interval.

Let $\alpha > 0$ and $\beta_1 \leq 0$ be arbitrary constants, let c_1 be the number of hits for a singleton window sequence and s the signature length. Then the window sequence is enlarged by

$$w \cdot \alpha \left(1 - \frac{c_1}{s}\beta_1\right) \tag{4}$$

nucleotides on both sides, where w is the window length. If $\beta_1 = 1$ and if we have a hit for each permutation, then $c1 = s$ holds and thus there is no enlargement.

Let β_2 be another arbitrary constant and c_1 and c_2 the number of hits for a window sequence of length 2. Let a and b be the start and end position of the window sequence (i.e. a is the first position of the left window and b is the last position of the right window). The resulting start position a' and end position b' of the interval are defined by:

$$a' \coloneqq a - n + q + \frac{c_2}{s} \cdot \beta_2 \cdot \frac{w - o + n - 2q}{2} \tag{5}$$

$$b' \coloneqq b + n - q - \frac{c_1}{s} \cdot \beta_2 \cdot \frac{w - o + n - 2q}{2} \tag{6}$$

If the number of hits are maximal and if $\beta_2 = 1$, then the resulting interval has the length n. The larger c_1 is, the more nucleotides of the read are probably located in the first window, so the right border of the interval can be contracted more than the left one.

If a window sequence consists of three or more windows, then the start and end position of the window sequence are used as interval. By default the parameters have the values $\alpha = 0.43$ and $\beta_1 = \beta_2 = 0.3$.

3.2 Variant tolerant alignment

Each interval is processed by a variant tolerant aligner. The aligner is based on Ukkonen's algorithm [17] and uses dynamic programming to calculate the optimal alignment according to the edit distance considering SNPs and indel variants.

Let $r = r_1, \ldots, r_m$ be the read and $t = t_1, \ldots, t_n$ be an interval on the reference genome. To handle SNPs the characters of reference genome are not elements of the set $\Sigma = \{A, C, G, T\}$, but of its power set $\mathcal{P}(\Sigma)$. For example, $t_i = \{A, C\}$ means that there is a SNP variant at position i and both nucleotides A and C co-exist at this position.

Let F be a matrix of size $m \times n$. The element $F[i, j]$ denotes the edit distance between $r_1 r_2 \cdots r_i$ and $t_{j'} t_{j'+1} \cdots t_j$ with some $j' \leq j$, so we have

$$F[i, 0] = i \quad 0 \leq i \leq m$$
$$F[0, j] = 0 \quad 1 \leq j \leq n$$

The other matrix elements are calculated using the following recursive formula if there are no indel variants.

$$F[i, j] = \min \begin{cases} F[i-1, j-1] + [\![p_i \notin t_j]\!] \\ F[i-1, j] \quad\quad +1 \\ F[i, j-1] \quad\quad +1 \end{cases} \tag{7}$$

Deletions can be considered by jumping back in the matrix. If there is a deletion variant of length k that ends at position j, such that the bases $t_{j-k}, \ldots, t_{j-1}$ are skipped, than the recursive formula for the column j is

$$F[i,j] = \min \begin{cases} F[i-1,j-1] & + [\![p_i \notin t_j]\!] \\ F[i-1,j] & +1 \\ F[i,j-1] & +1 \\ F[i-1,j-k-1] & + [\![p_i \notin t_j]\!] \\ F[i,j-k-1] & +1 \end{cases} \tag{8}$$

If there are more deletions that ends at positions j, they can be added to the recursion formula by appending two more terms to the minimum expression. So the recursion to be used at each reference position j is determined by the number and length(s) of the deletion variants ending at position j.

Handling insertions is slightly more complex. If there is an insertion $s_1 s_2 \cdots s_k$ before position j (such that a string containing this insertion is for example $t_{j-1} s_1 \cdots s_k t_j$), then we need k optional extra columns in the matrix F. To calculate the column for s_1 we have to access the column for t_{j-1}. The columns $s_2, \cdots, s_k$ can be determined straight-forwardly using the above recursion (7). To calculate the column t_j we not only have to access the column t_{j-1}, but also s_k (analogous to equation (8)). So for each insertion that ends before position j two terms are added to the minimum expression.

In contrast to the index of VATRAM, the aligner uses all available variants.

Filling the whole matrix consumes too much time. Therefore the user can set an error threshold k, so that only the parts of the matrix whose values are less or equal to this threshold are calculated. The details of this technique are more complicated, since there can be insertion or deletion variants and you have to ensure that all matrix elements that are accessed are already calculated.

The basic idea is to store for each matrix element an additional information called *next row* which indicates the row index of the next element in the given column that is not larger than k. The *next row* information of column j is written during the calculation of the column $j-1$ in F. Before an uninitialized element in F will be accessed, it is set to $k+1$ beforehand. It can be shown that it does not matter that the correct value of that element is maybe larger than $k+1$. In [9] the details of this acceleration method are extensively described and the correctness of this pruning rule is proven.

The alignment is done for each given interval. If the user is only interested in the best mapping, then the alignment cost of the first alignment can be used as error threshold for the second alignment to speed up the alignment process.

3.3 Paired-end reads and long reads

VATRAM is also able to map paired-end reads. For that purpose both read sequences u and v of the paired-end read are first processed separately. After creating the list of window sequences L_u and L_v for both read sequences, we look for window sequence pairs $(l_u, l_v) \in L_u \times L_v$, such that the distance between l_u

and l_v fits to the insert size distribution of the paired-end reads. To rank the window sequence pairs, the scoring values are combined to a single value. Finally, the alignment is done separately for both read sequences.

So far, we assume that all reads have a constant (or nearly constant) length n. However, some sequencing machines produce reads which length varies considerably. For example PacBio reads [6] have a length from a few hundred nucleotides up to 60 000. The window based approach of VATRAM needs reads of more or less constant length. Thus, variably long reads are split into fragments whose length $\tilde{n}$ fits to the window length w. Then analogous to paired-end reads for each fragment the corresponding window sequence lists $L_1, \ldots, L_k$ are calculated. Afterwards, we look for tuples $(l_1, \ldots, l_k) \in L_1 \times \cdots \times L_k$ of windows sequences whose distances fits to the distance of the read fragments. The scores of the window sequences are combined to a single value which is used for ranking the window sequence tuples. If some fragments at the end or maybe in the mid of the read are not found, then there are some gaps in the tuple, so we allow that each l_i is equal to an empty window sequence (or more formally $(l_1, \ldots, l_k) \in L'_1 \times \cdots \times L'_k$ where $L'_i := L_i \cup \{\emptyset\}$ for all $1 \leq i \leq k$). In this case the tuple gets a lower score, but the read can still be mapped if the tuple represents the correct position. After ranking, each tuple is converted into an interval which is processed by the aligner. The splitting procedure is applied if the fraction n/w is larger than a given constant f (by default $f = 0.8$ is used).

4 Experiments

We compare our read mapper VATRAM to other read mappers. These are BWA-SW [13], as Bowtie2 [10], BWA-MEM [11] in standard configuration and BWA-MEM configured such that it uses the edit distance as metric for the alignment (denoted as "BWA-MEM in edit-distance configuraiton"). Moreover we have tested *mrsFastUltra* [8], which is a variant tolerant read mapper like VATRAM. Our read mapper was executed with and without using the known variants to quantify the benefit of using variants.

VATRAMs parameter configuration is shown in table 1. The formula for the signature length is an empirical result that leads to an approximately constant memory usage of about 14 GB independent of the window distance respectively read length. For $n = 100$ the signature length is $s = 36$.

4.1 Simulated reads

In our first experiment we used simulated single-end reads of constant length to compare VATRAM with other read mappers. Using simulated reads has the advantage over real reads that the correct position is known.

The reads were created from the human reference genome (GRCh37). Each known variant is inserted into the reads with its population frequency. Furthermore sequencing errors respectively unknown variants are inserted in the reads

Table 1. Parameters used for the experiments. The splitting procedure was only used for data sets whose reads have variable length. In this case the parameters are $w = 140$, $o = 125$ and $s = 36$. The lower part of the table shows parameters that are only needed for read mapping and not for index creation (in contrast to the parameters in the upper part).

Symbol	Meaning	Value
w	window length	1.4n
o	window distance	1.25n
q	q-gram length	17
s	signature length	$\left\lceil \frac{8.63n}{22.6+0.0138n} \right\rceil$
δ	variant consideration threshold	20%
l	variant combination limit	3
κ	maximal number of selected window sequences	64
α	interval calculation	0.43
β_1	interval calculation	0.3
β_2	interval calculation	0.3
$\tilde{n}$	fragment length for reads with variable length	100
f	a read is split if $n/w > f$	0.8

with a probability of 2%, 4% or 8% per position. Of the errors, 90% are substitutions of single nucleotides, the remaining 10% are insertions or deletions of variable length. The ratio of the correct and wrong mapped reads as well as the runtime and the memory consumption are shown in figure 3 for different read lengths. Note that the sum of the correct and wrong mapped reads is not equal to 100%, because there can be reads that remain unmapped by the particular read mapped, so the sum is equal to the fraction of the reads that were mapped.

As we can see the memory consumption of VATRAM (about 14 GB) is significantly higher than those of the other read mappers. As most modern workstations have 16 GB of memory or even more, we do not expect this to be a major disadvantage.

The runtime of VATRAM is (in most cases) higher than the runtime of BWA or Bowtie2. The reason for this is that VATRAM takes variants into account and thus solves a problem that is more complex than the alignment that is done by BWA or Bowtie2. In fact, most of the runtime of VATRAM is needed to calculate the variant tolerant alignment. The mapping process (i.e. calculating the signature values of a read, looking them up to get the window indices and converting them to intervals) is much faster (approximately five times) than the alignment calculation[16]. The read mapper *mrsFastUltra* needs more time than VATRAM if the error rate is less or equal to 4%. With an error rate of 8%, *mrsFastUltra* is faster than VATRAM, but then over 50% of the reads are left unmapped.

VATRAM maps more reads to the correct position than *mrsFastUltra* which is also variant tolerant. However, BWA-MEM produces even better results (especially for the short reads) although it has no information about the known variants. The number of wrongly mapped reads of VATRAM is small, similar

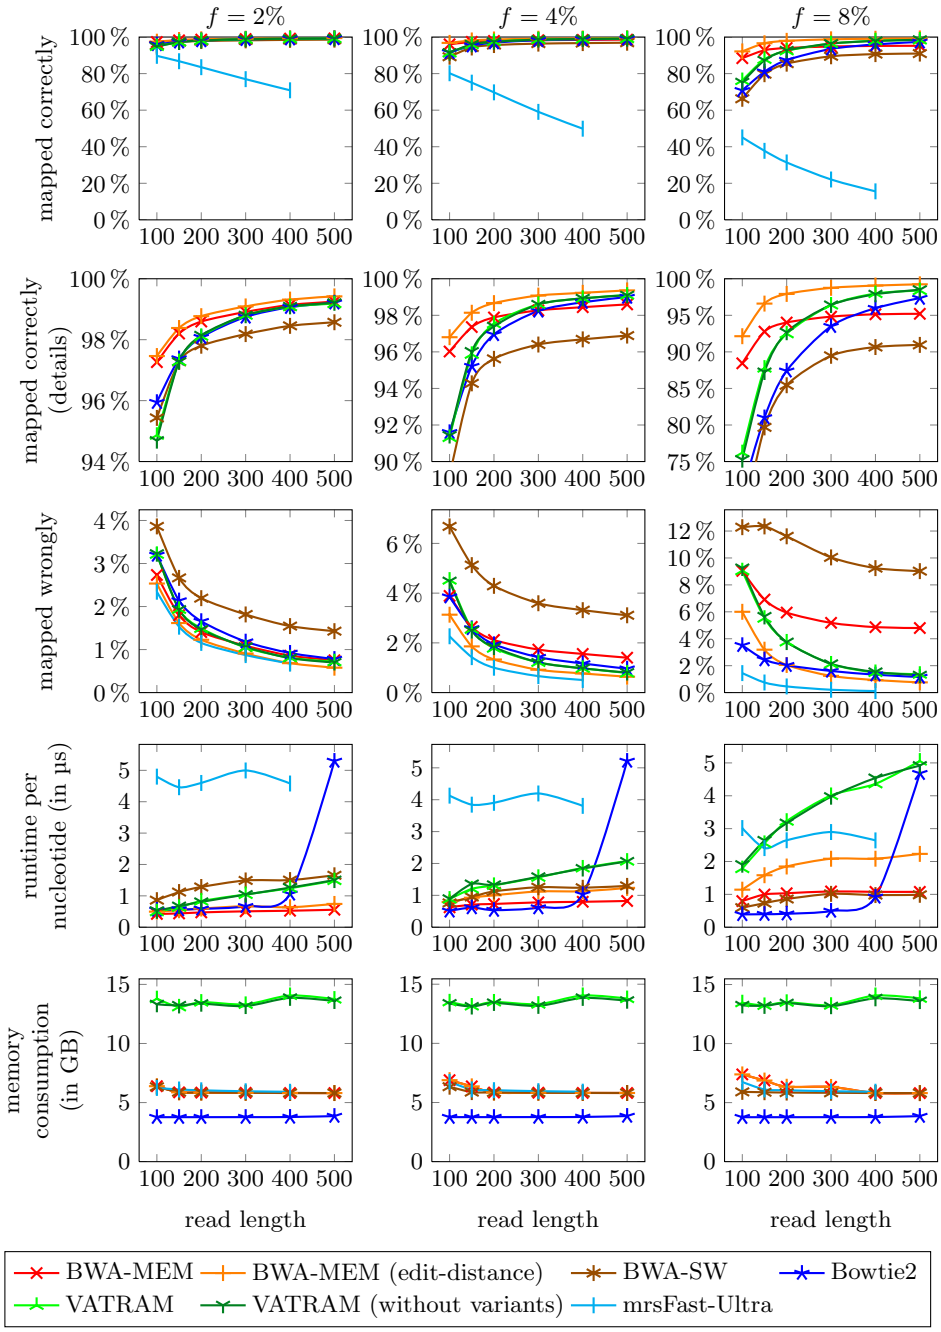

Fig. 3. Comparison between different readmappers using simulated reads. The error rate was 2%, 4% and 8% per position (the error rate in the diagrams of one column is constant and shown above the diagrams of the first row). The x-axis shows the read length. The ratio of correctly mapped reads is shown twice (first and second row) to improve the visibility.

to BWA-MEM and Bowtie2. The lowest fraction of wrongly mapped reads is achieved by *mrsFastUltra*. However, this is not surprising, since *mrsFastUltra* can only map a small number of reads, so many reads that are difficult to map correctly are left unmapped.

Although the simulated reads contain variants, there is no great benefit to consider them with VATRAM. Only for a read length of $n = 100$ and an error of 2%, the improvement is visible in figure 3. However, the fraction of correctly mapped reads only increases by 0.18%-points to 94.86%. This is considerably fewer than the results of BWA-MEM in edit-distance configuration with 97.4%. The reason for this small improvement is that most known variants appears with a frequency of less than 20%. These variants are not considered by VATRAMs index, because the variant consideration threshold was $\delta = 0.2$. A lower threshold would lead to even worse results [16], since reads that do not contain the variant (these are more than 80%) are found with a lower probability if the variant is added to the index.

In this experiment, VATRAM did not produce better results than BWA-MEM, although it uses extra knowledge about the known variants. However, VATRAM outperformed the variant-toleratn read mapper *mrsFastUltra*. The benefit of considering the variants was very low.

4.2 Real reads

In this section we use real data sets to compare the read mappers. When using real reads, the correct position is not known. One solution is to compare only the number of mapped reads. However, this is not a reasonable measure, because one may map a read to an arbitrary position which would lead to the best possible performance under this evaluation metric. Another method is to compare the edit distance of the different alignments. However, calculating the edit distance between a read and the corresponding interval on the reference genome does not consider any variants and therefore the unknown correct position of the read would not necessarily have the smallest distance.

Therefore we decided to use VATRAM's aligner to realign the reads to the interval on the reference genome given by the different read mappers. A mapping is defined as correct if there is no other mapping that leads to a lower edit distance considering all known variants. If two or more different mappings have the same optimal alignment cost, then all are treated as correct. Note that this does not provide VATRAM with an advantage because the aligner of VATRAM is just a tool to calculate the optimal alignment considering all known variants.

Table 2 shows an overview of the different data sets. There are two paired-end data sets with a constant read length as well as three single-end data sets with a variable read length created with three different sequencing machines.

Real paired-end reads In Figure 4 the results of the two paired-end data sets are shown. There are no values for *mrsFastUltra*, because the program crashes for unknown reason.

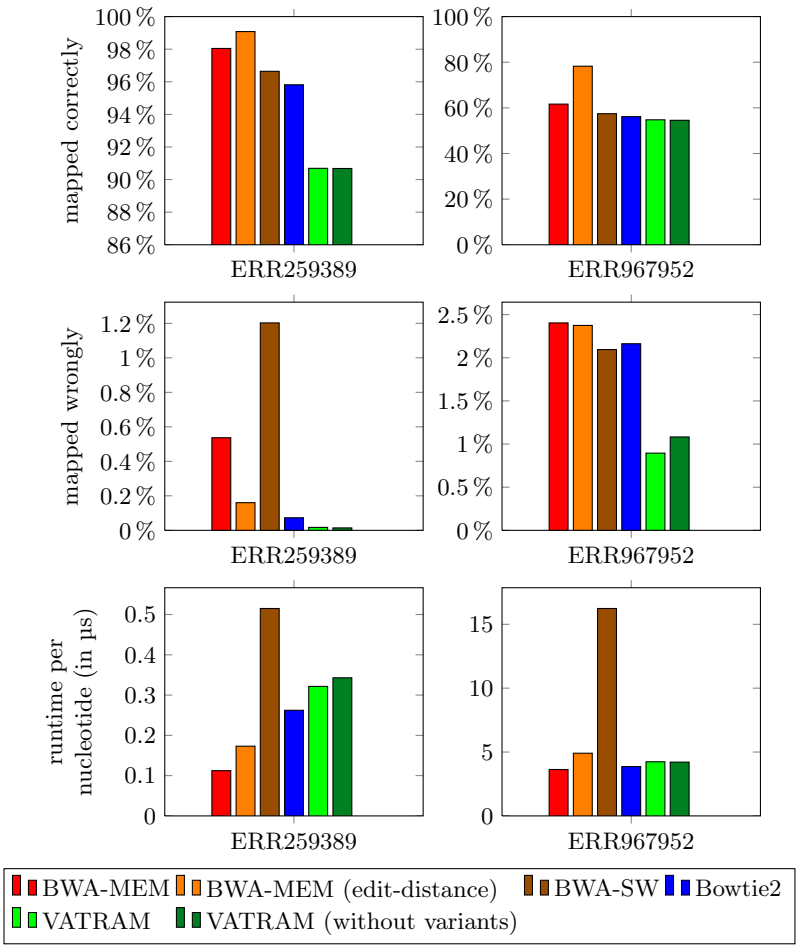

Fig. 4. Real datasets with paired-end reads (see Table 2). *mrsFastUltra* crashes on both datasets.

Table 2. Data sets used in the experiments. The data set name refers to NCBI sequence read archive (SRA). Column n shows the average read length, column *Type* indicates whether the data set contains paired-end (PE) ord single-end (SE) reads and the column *Const?* shows if all reads in the data set have the same length (*yes*) or not (*no*). The last column shows the error threshold that was used for the read mapper VATRAM. Note that the later alignment that is executed for each read mapper is done with a much higher error threshold to ensure that always the best alignment is calculated.

Name	Sequencing machine	n	Type	Const?	Err.thr.
ERR259389	Illumina MiSeq	151	PE	yes	15
ERR967952	Illumina MiSeq	250	PE	yes	25
DRR003760	Illumina MiSeq	180	SE	no	20
SRR003174	454 GS FLX Titanium	565	SE	no	60
SRX533609	PacBio RS II	8651	SE	no	450

For the first dataset (*ERR259389*) VATRAM achieves with 0.017% an extremely low rate of wrongly mapped reads. The number of reads that Bowtie2 mapped wrongly are more than 4 times larger and the other read mappers produce even worse results according to this measure. On the other hand, only 90.7% of the reads are mapped correctly by VATRAM. The other read mappers achieve 95.8% or more; so many reads remain unmapped when using VATRAM.

In the dataset *ERR967952* the ratios of correctly mapped reads are almost equal for VATRAM, Bowtie2, BWA-SW and BWA-MEM in standard configuration. However, BWA-MEM in edit-distance configuration is able to map about 17%-points more reads correctly than the other read mappers. On the other hand the number of wrongly mapped reads of BWA-MEM (in both configuraitons), Bowtie2 and BWA-SW is always twice as large as the number of reads that VATRAM mapped wrongly. Which read mapper is better depends on the application: If you want to have many correctly mapped reads, it makes sense to use BWA-MEM in edit-distance configuration, if you want to minimize the number of wrongly mapped reads, it is better to use VATRAM. The runtime of VATRAM in this dataset is comparable with the runtime of Bowtie2 and BWA-MEM. This is interesting, because the alignment process of VATRAM is usually much more time consuming due to the consideration of variants.

For both data sets there is only a slight benefit when the variants are added to the index. According to the number of correctly mapped reads, there is no change visible in Figure 4. The ratio of wrongly mapped reads decreases for the data set *ERR967952* by about 17%. The runtime was reduced by approximately 6% in the dataset *ERR259389*. For the other data sets there was no significant change.

Real reads with variable length So far we only analyzed data sets where all reads have the same length. However, sequencing machines like “454 GS” FLX or PacBio produce reads with varying length. In Figure 5 the results for three such data sets are shown.

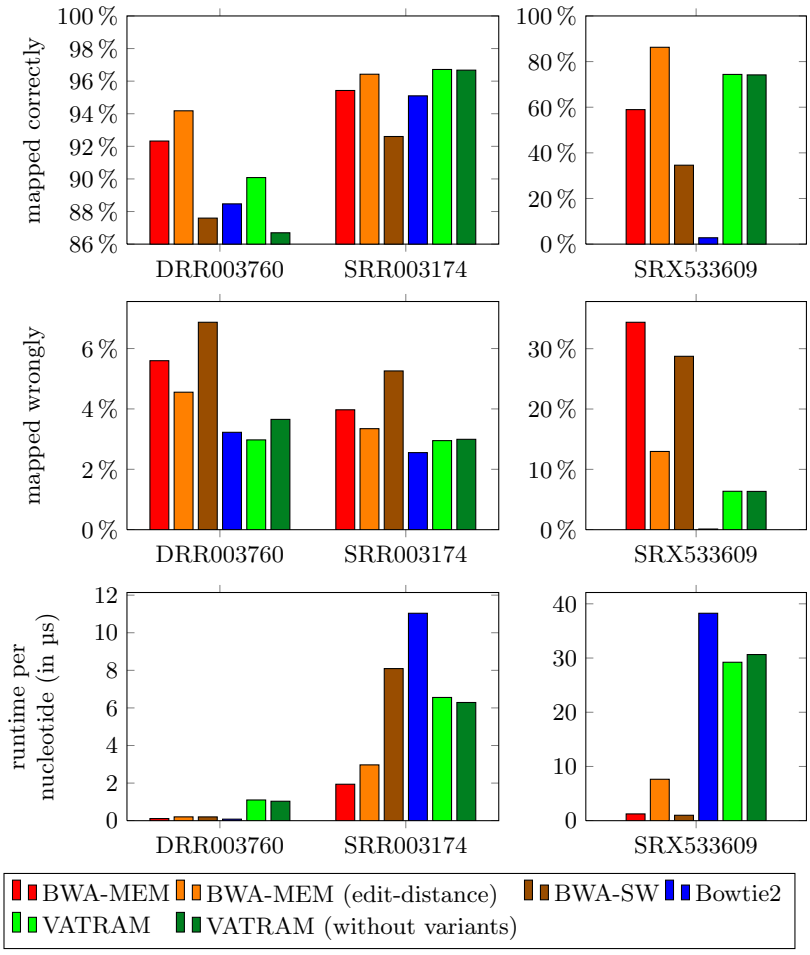

Fig. 5. Real reads with variable length. Information about the data sets can be found in Table 2. There are no results for *mrsFast-Ultra*, because it is not able to map reads with variable length.

The data set *DRR003760* is interesting, because considering the variants leads to clearly better results for VATRAM. The ratio of correctly mapped reads increases from 86.7% to 90.1%. Such a high improvement of 3.4%-points was not measured in other datasets, neither for the synthetic nor for the real data sets. The reason for this is that the reads of the *DRR003760* are located in the human HLA region which contains much variants in a very small area. Therefore, there is a great benefit, if the known variants are added to VATRAMs index. However, the ratio of correctly mapped reads is still lower than the corresponding ratio of BWA-MEM (92.3% in standard respectively 94.2% in edit-distance configuration).

The data set *SRR003174* shows that VATRAM is able to produce better results than BWA-MEM: Not only is VATRAMs ratio of wrongly mapped reads smaller than the ratio of BWA-MEM (in both configurations), but VATRAM also maps more reads to the correct position, even for the edit distance configuration of BWA-MEM which always leads to the highest number of correctly mapped read in the other experiments.

The data sets *DRR003760* and *SRR003174* contain reads with an moderate read length of 180 respectively 565 nucleotides on average. The last column in figure 5 shows the results for a read data set generated by a PacBio sequencing machine whose reads are much longer (*SRX533609*). VATRAMs aligner was not constructed to align such long reads with 10 000 nucleotides or even more. Therefore its running time is more than 20 times longer than the running of BWA-MEM in standard configuration. Only Bowtie2 needs even more time than VATRAM, since it was not designed for such long reads either. According to the number of correctly and wrongly mapped reads VATRAM performs clearly better than BWA-MEM in standard configuration. However, if the edit distance configuration of BWA-MEM is used (which makes more sense, since the edit distance is used as measure to determine whether a read is mapped correctly), then the ratio of correctly mapped reads is 12%-points better than VATRAMs ratio. However, also the ratio of wrongly mapped reads is 6.6%-points greater than the corresponding ratio of VATRAM. Which read mapper is better in this case depends on the application. VATRAM achieves the best precision, however the recall of BWA-MEM in edit-distance configuration is better than VATRAMs recall.

5 Discussion and conclusion

In the last decades many read mappers were developed, so there is the question why one should develop another read mapper. Most of the commonly used read mappers do not use knowledge about known variants. However, there are many differences between the reference genome and the genome of one individual. By now about 150 millions of variants are known[7] whose usage may improve the mapping process significantly. Common read mappers (like BWA or Bowtie2)

[7] `https://www.ncbi.nlm.nih.gov/dbvar/content/org_summary`, access date: 20th January 2017.

treat these variants similarly to sequencing errors and thus may have problems to map reads containing known variants to the correct position. Therefore we developed a new read mapper called VATRAM based on min-hasing that is able to consider known variants.

In section 4 we compared VATRAM with other read mappers. In most cases BWA-MEM in edit-distance configuration performs best according to the ratio of correctly mapped reads. For the data set *SRR003174*, VATRAM was the best read mapper. An additional strength of VATRAM is the very low ratio of wrongly mapped reads. It was usually much better than the ratio of Bowtie2 or BWA.

We tested not only common, non-variant tolerant read mappers like BWA or Bowtie2, but also the variant tolerant read mapper *mrsFastUltra*. However, its results were not convincing according to the very low rate of correctly mapped reads. Furthermore, it failed to handle all five tested real data sets.

VATRAM needs more time than BWA or Bowtie2 to process a data set. This is not surprising, since calculating a variant tolerant alignment is more complex and therefore more time consuming than a usual alignment that is for example performed by Bowtie2 and BWA.

We executed VATRAM with and without the information about the known variants. The benefit of considering the variants is usually small, because only a small ratio of the reads contain variants. However, for the data set *DRR003760* which contains reads from a region with many variants in a small area the mapping quality of VATRAM improved significantly after adding the known variants to the index.

All in all, VATRAM is a competitive read mapper in comparison to BWA and Bowtie2. Currently VATRAM is implemented as a prototype, and certainly further optimizations in terms of speed and memory usage are possible. An open question is in how many individuals a variant must appear such that adding this variant to the index leads to better results. This question could also be analyzed theoretically (using lemma 1), independent from the details of VATRAMs implementation.

References

1. Berlin, K., Koren, S., Chin, C.S., Drake, J.P., Landolin, J.M., Phillippy, A.M.: Assembling large genomes with single-molecule sequencing and locality-sensitive hashing. Nature biotechnology 33(6), 623–630 (2015)
2. Broder, A.Z.: On the resemblance and containment of documents. In: Compression and Complexity of Sequences (SEQUENCES'97). pp. 21–29. IEEE (1997)
3. Broder, A.Z., Charikar, M., Frieze, A.M., Mitzenmacher, M.: Min-wise independent permutations. In: Proceedings of the 30th annual ACM symposium on Theory of computing (STOC). pp. 327–336. ACM (1998)
4. Buhler, J.: Efficient large-scale sequence comparison by locality-sensitive hashing. Bioinformatics 17(5), 419–428 (2001)
5. Casperson, M.: Minhash for dummies. `http://matthewcasperson.blogspot.de/2013/11/minhash-for-dummies.html` (November 2013)

6. Eid, J., Fehr, A., Gray, J., Luong, K., Lyle, J., Otto, G., Peluso, P., Rank, D., Baybayan, P., Bettman, B., Bibillo, A., Bjornson, K., Chaudhuri, B., Christians, F., Cicero, R., Clark, S., Dalal, R., deWinter, A., Dixon, J., Foquet, M., Gaertner, A., Hardenbol, P., Heiner, C., Hester, K., Holden, D., Kearns, G., Kong, X., Kuse, R., Lacroix, Y., Lin, S., Lundquist, P., Ma, C., Marks, P., Maxham, M., Murphy, D., Park, I., Pham, T., Phillips, M., Roy, J., Sebra, R., Shen, G., Sorenson, J., Tomaney, A., Travers, K., Trulson, M., Vieceli, J., Wegener, J., Wu, D., Yang, A., Zaccarin, D., Zhao, P., Zhong, F., Korlach, J., Turner, S.: Real-time DNA sequencing from single polymerase molecules. Science 323(5910), 133–138 (2009)
7. Ferragina, P., Manzini, G.: Indexing compressed text. Journal of the ACM 52(4), 552–581 (2005)
8. Hach, F., Sarrafi, I., Hormozdiari, F., Alkan, C., Eichler, E.E., Sahinalp, S.C.: mrsFAST-Ultra: a compact, SNP-aware mapper for high performance sequencing applications. Nucleic Acids Research 42(Webserver-Issue), 494–500 (2014)
9. Kramer, B., Quedenfeld, J., Schrinner, S., Bargull, M., Benadjemia, K., Stricker, J., Losch, D.: VATRAM – VAriant Tolerant ReAd Mapper. Tech. rep., Project Group PG583, Computer Science, TU Dortmund, Germany (2015)
10. Langmead, B., Salzberg, S.L.: Fast gapped-read alignment with Bowtie 2. Nature methods 9(4), 357–359 (2012)
11. Li, H.: Aligning sequence reads, clone sequences and assembly contigs with BWA-MEM. arXiv preprint arXiv:1303.3997 (2013)
12. Li, H., Durbin, R.: Fast and accurate short read alignment with Burrows–Wheeler transform. Bioinformatics 25(14), 1754–1760 (2009)
13. Li, H., Durbin, R.: Fast and accurate long-read alignment with Burrows-Wheeler transform. Bioinformatics 26(5), 589–595 (2010)
14. Metzker, M.L.: Sequencing technologies—the next generation. Nature Reviews Genetics 11(1), 31–46 (2010)
15. Navarro, G.: Compact Data Structures: A Practical Approach. Cambridge University Press, Cambridge (009 2016)
16. Quedenfeld, J.: Variantentolerantes Readmapping durch Locality Sensitive Hashing. Master's thesis, Computer Science XI, TU Dortmund, Germany (2016)
17. Ukkonen, E.: Finding approximate patterns in strings. Journal of Algorithms 6(1), 132 – 137 (1985)